

**Energy Efficient High Performance Computing Working Group
04/09/2019 Meeting Report**

INTRODUCTION

The Energy Efficient High Performance Computing Working Group (EE HPC WG) held a meeting on 4/09/19. This Working Group is composed of members representing major governmental departments and independent agencies, private sector representatives, and members of the academic community. More information can be found at the working group's website, <https://eehpcwg.llnl.gov/>.

NEXT MEETING: Tuesday, June 11th, 2019 9:00-10:00AM Pacific Time

ANNOUNCEMENTS:

Siddhartha Jana announced that there will be a One-day Symposium on "Energy Efficiency Challenges for High Performance Computing (HPC) Systems" on May 17, 2019 at C-DAC, KP Bengaluru. Both Natalie and Sid will be 'virtually' presenting.

Siddhartha also announced that the second annual face-to-face Power Stack meeting is coming up. It will be held the week prior to ISC2019 and is taking place at LRZ. For more information, please see: <https://www.hipeac.net/news/6895/powerstack-a-global-response-to-the-power-management-problem-for-exascale/>

Andrea Bartolini announced that there will be a Special Session on Tools for Energy Efficient Computing in Bialystok, Poland from September 8- 11. For more information, see: https://www.ppam.pl/docs/workshops/PPAM_2017_Special_Session_on_EE.pdf

Andrea also announced that the HiPEAC conference will be held in Bologna, Italy January 20-22, 2020. For more information, please see: <https://www.hipeac.net/2020/bologna/#/>

GRID INTEGRATION TEAM REPORT:

Grant Stewart reported for the Grid Integration Team. We're working on our current paper, which is the fourth in a series for this team. This paper is about grid integration of growing super computers with highly dynamic power flow requirements. We've completed questionnaires and interviews with four sites. We're analyzing the data that we've collected. We've started to create a table that sums up what the grid conditions are for each site and how they're preparing for this this growing dynamic demand. The sites are ORNL, LLNL, NCAR and LANL. We're hoping to have a four page paper to submit to the EE HPC SOP workshop.

LIQUID COOLED RACK SPECIFICATION TEAM REPORT:

David Sickinger reported for the Liquid Cooled Rack Specification Team. This Team will work with Open Compute Project (OCP), ASHRAE and perhaps others to develop and disseminate a liquid cooled rack specification. The Team has identified approximately nine subtopics for an open specification of a liquid-cooled rack. This specification does not include the primary water or the CDU secondary loop. The specification does cover the fluid and hardware inside the rack and up to the point where hardware connects to the incoming fluid. It includes the water connectors to that incoming fluid. The team has initiated two subgroups to focus on topics that are closely linked. The first one is the transfer fluid and the second is the wetted materials requirements. Both groups have had a kickoff meeting as well as a follow-up meeting.

The wetted material subgroup is working with the OCP material compatibility list as a starting point. This version of the list is very short to ensure compatibility and reduce risk in terms of the transfer fluid subgroup. The wetted material subgroup decided upon a multi-vendor non-glycol water-based spec as the starting point. The OCP is starting off with a 25% propylene glycol solution from two different vendors, but this subgroup decided to focus on water-based solutions.

The transfer fluid subgroup has recently been discussing the frequency of fluid testing and different folks are thinking about whether or not a graded approach could be taken. Frequency of testing could be left to the owner and operators with an understanding of greater risk, especially relative to the bio contamination, with less frequent testing. There is a trade-off of testing versus how much time you have to spend with these maintaining these systems. The subgroup is attempting to collect data from water-cooling users and suppliers in order to determine the appropriate test process as well as test intervals.

POWERAPI TEAM REPORT:

Ryan Grant reported for the PowerAPI Team. The Power API is a holistic specification for power monitoring and control on large-scale systems all the way down to small scale systems. It's basically an abstract device interface so that you can request basic things- like power and energy- to find out more about your system. It has a basic hierarchy that lets you navigate through all the objects in a system. Objects can be anything from the entire chassis, rack nodes and all the way down to individual cores. You can access power and energy consumption and even get things like temperatures. The specification exists and is available on the internet. We have a GitHub page and there's a link to it from the EE HPC WG website.

Currently, we're preparing for all of the interactions with Power Stack. We have been working on some text for integration with some of the things that Power Stack is trying to do. We're looking a lot at our job management interfaces and how we manage jobs as a whole as well as be able to pass around general information within a system that relate to all the things that we need for job management. We are looking at some reliability aspect to it too. We're also iterating on some proposals for new reporting types that come out of the system so we can also create higher level reports and summaries of system operation current status. We're looking at adding in some new

capabilities having to do with reporting, both improving the efficiency as well as adding new features.

OPERATIONAL DATA ANALYTICS TEAM REPORT

Natalie Bates reported for the Operational Data Analytics (ODA) Team. This is a team that is focused on sites that have extensive instrumentation, data collection and data analysis tools. It is looking at experiences and best practices that are associated with these kinds of capabilities. We have two thrusts in this activity; one is a global survey of sites that have actually implemented and/or are implementing these kinds of instrumentation data collection and data analysis systems that span both the infrastructure and the HPC system. We have identified some of these sites and are in the process of collecting information from them to understand how they're using their ODA capabilities.

In addition to the global site survey, we have been working directly with Lawrence Berkeley National Lab on some case studies. They have a pretty sophisticated system that's being used on an ongoing basis for their entire facility and all of their HPC systems. We're hoping to submit two of these case studies to the EE HPC SOP Workshop.

ENERGY AND POWER AWARE JOB SCHEDULING AND RESOURCE MANAGEMENT (EPAJSRM) TEAM REPORT

Natalie Bates reported for the EPAJSRM Team. We've written and published three papers from this team; two to workshops and one to a journal. These papers described the results of a global site survey for those sites who had implemented or were planning to implement EPAJSRM on a production large-scale system.

The Team is currently looking at site policies and thinking about optimization techniques for policies, which would implement energy and power aware capabilities. It's a small subgroup. We're looking at optimization work already completed for other non-data center facilities.

Sid reported that the Power Stack Initiative is also trying to figure out how some of the high-level facility plant level policies or decisions with regards to: power capping a system, or ramp up ramp down rates, or interaction with external electric grids. We want to understand how some of those high-level decisions trickle down to system level operations. How do you even translate this high-level policy decisions to an entity like a job scheduler or a workload manager?

PROCUREMENT CONSIDERATIONS TEAM REPORT

Jim Laros and Steve Martin reported for the Procurement Considerations Team. Jim reported that the Team has been restructuring the document and taking a taking a critical look at it to serve the community now as opposed to how we targeted serving the community previously. One of the original sections on measurement, for example, is probably too detailed at this point. I mean at first it was necessary to be that instructive, I think now the state of the art has certainly caught up and we need to revisit that section. I think our plan is to retain that

content because it's probably useful for people, but it's not going to be in this version of the document.

Steve then reported that he was at the last meeting and it was actually a little bit light on attendance. Everybody had been given some homework to respond to the suggested changes in the formatting and in the wording in the introduction. Feedback was given at that meeting, although some other people had sent in their feedback prior to the meeting. The new format is in pretty good shape. There were a couple of minor comments on a table in the new section, but otherwise the format looks good. The introduction also lays the ground work out well and describes the scope of the document and the intent.

CONFERENCES REPORT

Siddhartha Jana reported on Conferences activities.

ISC2019 is scheduled for the week of June 16-20 in Frankfurt. The EE HPC WG has submitted a BoF for the Green500. The EE HPC WG has been collaborating with the Green500 and the Top500 to improve the methodology by which power is measured while doing a full-system benchmark run- like MPLinpack. Acceptance notice is tomorrow, so we'll know then the status of the submission. There are other activities at ISC2019 that should be of interest to the EE HPC WG. There is a full day Power Stack Workshop scheduled for Thursday, June 20th. There is a Power Stack project poster. Also, there is a half-day tutorial on GEOPM, which is an open source power management runtime and Intel will host a separate conference room to highlight other vendor plug-ins for GEOPM.

SC2019 is scheduled for the week of November 17-22 in Denver. The annual EE HPC WG workshop has been accepted and is tentatively scheduled for all day on Monday. right now. The Datacenter Automation, Analytics and Control (DAAC) Workshop was has an overlap in interest and audience with the EE HPC WG Workshop, but fortunately it is scheduled for Friday. The EE HPC WG booth has been secured. April 23rd is the deadline for panel submissions and the EE HPC WG is planning to submit one with a 'big data and artificial intelligence' focus. BoF and Poster submission deadlines are the end of July.

The EE HPC State of the Practice Workshop is planned as part of the International Conference for Parallel Processing (ICPP) Conference. It is scheduled for August 5th in Kyoto. This is the first-ever EE HPC WG workshop in the traditional sense with solicited papers from the community. Satoshi Matsuoka will be the keynote speaker.

PARTICIPANTS INCLUDED

Name	Organization
Natalie Bates	EE HPC WG
Siddhartha Jana	Intel
Grant Stewart	LANL
Xingfu Wu	Argonne NL
David Strickland	NREL
Ryan Grant	Sandia NL
Kate Berard	DOE
Ralph Wescott	PNNL
Jason Zeiler	Coolit Systems
Steve Martin	Cray
Bob Boltz	Aquila
Andrea Bartollini	UNIBO
James Laros	Sandia NL
Daniel Robinson	Formerly FEMP
Valerie Taylor	Argonne NL